

10/532474

PCT/AU03/01418



Patent Office
Canberra

REC'D 17 NOV 2003
WIPO PCT

I, JANENE PEISKER, TEAM LEADER EXAMINATION SUPPORT AND SALES hereby certify that annexed is a true copy of the Provisional specification in connection with Application No. 2002952274 for a patent by IDEADATA PTY LTD as filed on 24 October 2002.



WITNESS my hand this
Tenth day of November 2003

JANENE PEISKER
TEAM LEADER EXAMINATION
SUPPORT AND SALES

**PRIORITY
DOCUMENT**

SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

BEST AVAILABLE COPY

002 THU 16:18 FAX 81 2 99255998

GRIFFITH HACK

@004

AUSTRALIA

Patents Act 1990

PROVISIONAL SPECIFICATION

Applicant(s):

IDEADATA PTY LTD

Invention Title:

**A COMPUTING DEVICE AND METHOD FOR RECORDING DATA
EXCHANGED BETWEEN ELECTRONIC DEVICES**

The invention is described in the following statement:

**A COMPUTING DEVICE AND METHOD FOR RECORDING DATA EXCHANGED
BETWEEN ELECTRONIC DEVICES**

FIELD OF THE INVENTION

5

The present invention relates to a computing device and method for recording data exchanged between electronic devices, and is of particular but by no means exclusive application to recording data packets transferred over a communications network.

10

BACKGROUND OF THE INVENTION

Recording data exchanged between electronic devices is desirable for several reasons. For instance, in the situation where the data being recorded includes data packets being transferred over a communications network, the record can be used to provide network administrators with an insight into the characteristics of the packets being transferred over their network. One such characteristic that network administrators are commonly interested in is the network address from or to which packets emanate or are destined. The address information assists network administrators in identifying potential points of congestion in their network, and as such allows the network administrator to re-configure their network to better handle the congestion.

Existing tools for recording data exchanged between electronic devices commonly create a separate entry for each piece of data exchanged between the devices. In the above example of data packets transferred over a communications network, the record maintained by existing tools would contain a separate entry for each packet exchanged over the network. Unfortunately, creating a separate entry for each piece of information (packet) has the potential to generate a very large record.

30

35

- 3 -

SUMMARY OF THE INVENTION.

According to a first aspect of the present invention there is provided a computing device for recording data exchanged between electronic equipment, the data having segments each of which has an attribute, the computing device including processing means for:

receiving the data;

10 identifying the segments in the data as being unique or equivalent, the attribute of the unique segment being different to the attribute of any other segment in the data, and the attribute of the equivalent segment being the same as the attribute of at least one other segment in the data;

15 creating a first record for each unique segment, the first record being representative of the unique segment; and

20 creating a second record for each set of equivalent segments, the second record being representative of the set of equivalent segments and is associated with a count field indicating a number of segments in the set of equivalent segments.

Thus, unlike existing tools which create a separate entry for each segment of the data, the computing device of the present invention creates only a single record (the second record) for segments which are equivalent to each other. The fact that the second record represents multiple segments is indicated by the count field, which can be used to stipulate the number of other segments. Use of the second record thereby minimizes the amount of storage space required to record data exchanged between electronic devices.

35 Preferably, identifying the unique segment and/or equivalent segment includes comparing the attribute of a segment with the attribute of any other segment.

-4-

Alternatively, the unique segment and/or equivalent segment can be determined by comparing the attribute of the segment with the first record and/or the second record.

5

Preferably, the first record and the second record each include a tag field and a data field, wherein the tag field identifies the attribute, and the data field includes the attribute.

10

Preferably, the first record and the second record include a delimiter for delimiting the combination of the tag field and the data field.

15

Preferably, the first record and the second record are in a format which is readily interpreted by a human.

20

Even more preferably, the format includes ASCII.

Preferably, the attribute of each of the segments includes information contained therein.

25

Alternatively, the attribute of each of the segments can include a characteristic thereof including the size of a segment.

30

Preferably, receiving the data is carried out over a period of time.

35

Even more preferably, the period is 10 minutes.

40

Preferably, the first record and second record each include a timestamp field which can be used to determine the validity of the respective first record or second record.

5

Preferably, each of the segments includes a data packet.

Alternatively, each of the segments can include fixed length cells.

Even more preferably, the data packet is an Internet Protocol (IP) data packet.

According to a second aspect of the present invention, there is provided a method for recording data exchanged between electronic devices, the data having segments each of which has an attribute, the method including the steps of:

receiving the data;

identifying the segments in the data as being unique or equivalent, the attribute of the unique segment being different to the attribute of any other segment in the data, and the attribute of the equivalent segment being the same as the attribute of at least one other segment in the data;

creating a first record for each unique segment, the first record being representative of the unique segment; and

creating a second record for each set of equivalent segments, the second record being representative of the set of equivalent segments and is associated with a count field indicating a number of segments in the set of equivalent segments.

30

Thus, unlike existing tools which create a separate entry for each segment of the data, the method of the present invention creates only a single record (the second record) for segments which are equivalent to each other.

The fact that the second record represents multiple segments is indicated by the count field, which can be used to stipulate the number of equivalent segments. Use

- 6 -

of the second record thereby minimizes the amount of storage space required to record data exchanged between electronic devices.

5 Preferably, identifying the unique segment and/or equivalent segment includes comparing the attribute of a segment with the attribute of any other segment.

10 Alternatively, the unique segment and/or equivalent segment can be determined by comparing the attribute of the segment with the first record and/or the second record.

15 Preferably, the first record and the second record each include a tag field and a data field, wherein the tag field identifies the attribute, and the data field includes the attribute.

20 Preferably, the first record and the second record include a delimiter for delimiting the combination of the tag field and the data field.

25 Preferably, the first record and the second record are in a format which is readily interpreted by a human.

Even more preferably, the format includes ASCII.

30 Preferably, the attribute of each of the segments includes information contained therein.

35 Alternatively, the attribute of each of the segments can include a characteristic thereof including the size of the segment.

Preferably, the step of receiving the data is carried out over a period of time.

Even more preferably, the period is 10 minutes.

Preferably, the first record and the second record each include a timestamp field which can be used to determine the validity of the respective first record or second record.

Preferably, each of the segments includes a data packet.

10

Alternatively, each of the segments can include fixed length cells.

Even more preferably, the data packet is an Internet Protocol (IP) data packet.

BRIEF DESCRIPTION OF THE DRAWINGS

Notwithstanding any other embodiments which may fall within the scope of the present invention, a preferred embodiment of the present invention will now be described, by way of example only, with reference to the accompanying figures, in which:

25 Figure 1 illustrates a computer system which includes an apparatus according to the preferred embodiment of the present invention;

Figure 2 illustrates the data exchanged in the system illustrated in figure 1;

Figure 3 illustrates a record which can be created by the apparatus shown in figure 1; and

35 Figure 4 illustrates another record which can be created by the apparatus shown in figure 1.

THE PREFERRED EMBODIMENT OF THE PRESENT INVENTION

As shown in figure 1, a computer system 1 which includes a first electronic device 3 and a second electronic device 5 that are interconnected to each other via a communication network 7. The electronic devices 3 and 5 are in the form of computer equipment such as personal computers, whilst the communication network 7 is a TCP/IP based network such as the Internet. The apparatus 9 is connected to the communication network 7 via a network node such as a router (not illustrated), and is a suitable programmed computing device such as a computer loaded with appropriate software and/or hardware.

The data 11 exchanged between the electronic devices 3 and 5 is made up of segments 11a and 11b, which are in the form of TCP/IP packets. As illustrated in figure 2, the segments 11a and 11b include an information field 13 that contains information. In the case of a TCP/IP packet, the information field 13 includes packet header information such as source IP address, destination IP address, protocol number etc.

As the segments 11a and 11b are exchanged between the devices 3 and 5, they will pass via the router as they traverse the network 7. Upon being received by the router, the router copies the segments 11a and 11b and forwards the copies thereof to the apparatus 9.

Alternatively, rather than being connected to the network 7 via the router, the apparatus 9 could be connected directly to the network 7. In this configuration the apparatus 9 would be capable of 'sniffing' the segments 11a and 11b as they traverse the network.

The apparatus 9 is configured to collect the copies of the segments 11a and 11b over a period of time, which

- 9 -

is ten minutes. Once the ten minute period is over, the apparatus 9 will proceed to process the copies of the segments 11a and 11b to determine whether the data 11 contains a unique segment which has an attribute that is

5 different to the attribute of other segments of the data and/or equivalent segments that share the same attribute. In the preferred embodiment, the attribute is the information contained in the information field 13 are the same as each other.

10.

The apparatus 9 determines whether the data 11 contains the unique segment or equivalent segments by comparing the information contained in the information field 13 of one of the segments 11a with the information

15.

contained in the information field 13 of another of the segments 11b. If the comparison determines that the information is different, then the apparatus 9 will consider the segment 11a to be the unique segment. On the other hand, if the comparison determines that the

20.

information is the same, then the apparatus 9 will consider the segments 11a and 11b to be the equivalent segments.

As can be seen in figure 2, the segments 11a and 11b include several information fields 13a, 13b and 13c. For the purposes of deciding which of the information fields 13a, 13b and 13c is to be used in the comparison, the appropriate information field 13a, 13b or 13c can be selected by the user of the apparatus 9. To enable the user to select the information field 13a, 13b or 13c, the apparatus 9 provides a graphical user interface from which the user can enter the appropriate information field 13a, 13b or 13c. The user can also select more than one information field 13 to be compared. In this case, the apparatus 9 uses the multiple selected information fields 13 in the previously mentioned step of comparing the information fields 13.

- 10 -

Once the apparatus 9 has identified the unique segment and/or equivalent segments, the apparatus 9 proceeds to create a first record 15 (illustrated in figure 3) if it determines the data 11 contains the unique segment. The apparatus 9 will also create the second record 17 (illustrated in figure 4) if it determines the data 11 contains equivalent segments. The first record 15 is representative of the unique segment and includes a count field 19a indicating that the first record represents a single segment of the data 11. The second record 17 is representative of the equivalent segments and also includes a count field 19b. However, unlike the count field 19a of the first record 15, the count field 19b of the second record 17 indicates the number of segments which are equivalent to each other. The count fields 19a and 19b contain information in hexadecimal format.

The records 15 and 17 each include a tag field 21a and 21b and a data field 23a and 23b. The data field 23a and 23b contain the information in the information field 13 of the segments 11a or 11b identified as the unique segment or equivalent segments. The tag field 21a and 21b is used to identify the information field 13. The information contained in the data field 23a and 23b is in hexadecimal form, whilst the tag field 21a and 21b is two letters in the ASCII format. By way of example, the tag field 21a and 21b could contain the letters 'DI' to indicate that the information in the data field 23a and 23b is the destination address of the segments 11a and 11b. The destination address being the network address of either the first electronic device 3 or second electronic device 5.

Whilst figures 3 and 4 illustrate the records 15 and 17 as having only one tag field 21 and data fields 23, the records 15 and 17 can have multiple tag fields 21 and data

11

fields 23. This would occur when the user selects more than one information field 13 from the graphical user interface. Where the records 15 and 17 have multiple tag fields 21 and data fields 23, each set of tag field 21 and data field 23 is delimited from other tag fields 21 and data fields 23 by, for example, the pipe character; that is, "|".

The records 15 and 17 include a timestamp field 25 which can be used to determine whether the records 15 and 17 are valid. The timestamp field 25 includes the actually time the first record 15 or second record 17 where created. The apparatus 9 is configured to automatically delete records which, for example, are more than 1 hour old because the record 15 or 17 is no longer relevant (valid).

When initially created, the records 15 and 17 are stored in the random access memory or hard disk of the apparatus 9. However, once created and initialised, the apparatus 9 can output the records 15 and 17 to a more permanent and readily accessible storage system 27. The more permanent storage system 27 is an SQL database. Whilst figure 1 shows the storage system 27 as being external to the apparatus 9, it is envisaged that the storage system 27 could form part of the apparatus 9. Where the storage system 27 is external to the apparatus 9, the storage system 27 and apparatus 9 are connected by a suitable link which would allow the records 15 and 17 to be transferred, such as an Ethernet link.

The following is a formal description of the algorithm used for recording the exchange of the data 11:

```
35  INP_LIST           //list of valid segments
    HASH              //hash table
    for each INP      //is a row from INP_LIST
```

- 12 -

INP.KEYS //Fields extracted from INP
INP.COUNTERS //counter fields extracted from INP
R //row returned from lookup of HASH (INP.KEYS)
if no R available, make new R as follows:
5 R.KEYS = INP.KEYS
R.COUNTERS += INP.COUNTERS
R.TI = INP.TI //TI being timestamp
field
10 R.DU = INP.DU //DU duration of validity
else update existing R as follows:
R.COUNTERS+=INP.COUNTERS
R.DU = max(R.TI + R.DU, INP.TI + INP.DU)
R.TI = min(R.TI, INP.TI)
15 endif
R is inserted into HASH(R.KEYS)

Those skilled in the art will appreciate that the invention described herein is susceptible to variations and modifications other than those specifically described. It should be understood that the invention includes all such variations and modifications which fall within the spirit and scope of the invention.

13

THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. A computing device for recording data exchanged between electronic equipment, the data having segments each of which has an attribute, the computing device including processing means for:
 - receiving the data;
 - identifying the segments in the data as being unique or equivalent, the attribute of the unique segment being different to the attribute of any other segment in the data, and the attribute of the equivalent segment being the same as the attribute of at least one other segment in the data;
 - creating a first record for each unique segment, the first record being representative of the unique segment, and
 - creating a second record for each set of equivalent segments, the second record being representative of the set of equivalent segments and is associated with a count field indicating a number of segments in the set of equivalent segments.
2. The computing device as claimed in claim 1, wherein identifying the unique segment and/or equivalent segment includes comparing the attribute of a segment with the attribute of any other segment.
3. The computing device as claimed in claim 1 or 2, wherein the first record and the second record each include a tag field and a data field, wherein the tag field identifies the attribute, and the data field includes the attribute.
4. The computing device as claimed in claim 3, wherein the first record and the second record include a delimiter for delimiting the combination of the tag field and the data field.

14

5. The computing device as claimed in any one of
the preceding claims, wherein the first record and the
second record are in a format which is readily interpreted
5 by a human.

6. The computing device as claimed in claim 5,
wherein the format includes ASCII.

10 7. The computing device as claimed in any one of
the preceding claims, wherein the attribute of each of the
segments includes information contained therein.

15 8. The computing device as claimed in any one of
the preceding claims, wherein receiving the data is
carried out over a period of time.

20 9. The computing device as claimed in claim 8,
wherein the period is 10 minutes.

10. The computing device as claimed in any one of
the preceding claims, wherein the first record and the
second record each include a timestamp field which can be
used to determine the validity of the respective first
25 record or second record.

11. The apparatus as claimed in any one of the
preceding claims, wherein each of the segments includes a
data packet.

30 12. The apparatus as claimed in claim 11, wherein
the data packet is an Internet Protocol (IP) data packet.

13. A method for recording data exchanged between
35 electronic devices, the data having segments each of which
has an information field containing information, the
method including the steps of:

15

receiving the data;
identifying the segments in the data as being unique or equivalent, the attribute of the unique segment being different to the attribute of any other segment in the data, and the attribute of the equivalent segment being the same as the attribute of at least one other segment in the data;

creating a first record for each unique segment, the first record being representative of the unique segment, and

creating a second record for each set of equivalent segments, the second record being representative of the set of equivalent segments and is associated with a count field indicating a number of segments in the set of equivalent segments.

14. The method as claimed in claim 13, wherein the step of identifying the unique segment and/or equivalent segment includes comparing the attribute of a segment with the attribute of any other segment.

15. The method as claimed in claim 13 or 14, wherein the first record and the second record each include a tag field and a data field, wherein the tag field identifies the information field, and the data field includes the information in the information field.

16. The method as claimed in claim 15, wherein the first record and the second record include a delimiter for delimiting the combination of the tag field and the data field.

17. The method as claimed in any one of the preceding claims, wherein the first record and the second record are in a format which is readily interpreted by a human.

- 16 -

18. The method as claimed in claim 17, wherein the format includes ASCII.

19. The method as claimed in any one of the preceding claims, wherein the attribute of each of the segments includes information contained therein.

20. The method as claimed in any one of the preceding claims, wherein the step of receiving the data is carried out over a period of time.

21. The method as claimed in claim 20, wherein the period is 10 minutes.

22. The method as claimed in any one of the preceding claims, wherein the first record and the second record each include a timestamp field which can be used to determine the validity of the respective first record or second record.

23. The method as claimed in any one of the preceding claims, wherein each of the segments includes a data packet.

24. The method as claimed in claim 23, wherein the data packet is an Internet Protocol (IP) data packet.

25. An apparatus substantially as herein described with reference to the accompanying figures.

26. A method substantially as herein described with reference to the accompanying figures.

Dated this 24th day of October 2002

35 IDEADATA PTY LTD

By their Patent Attorneys

GRIFFITH HACK

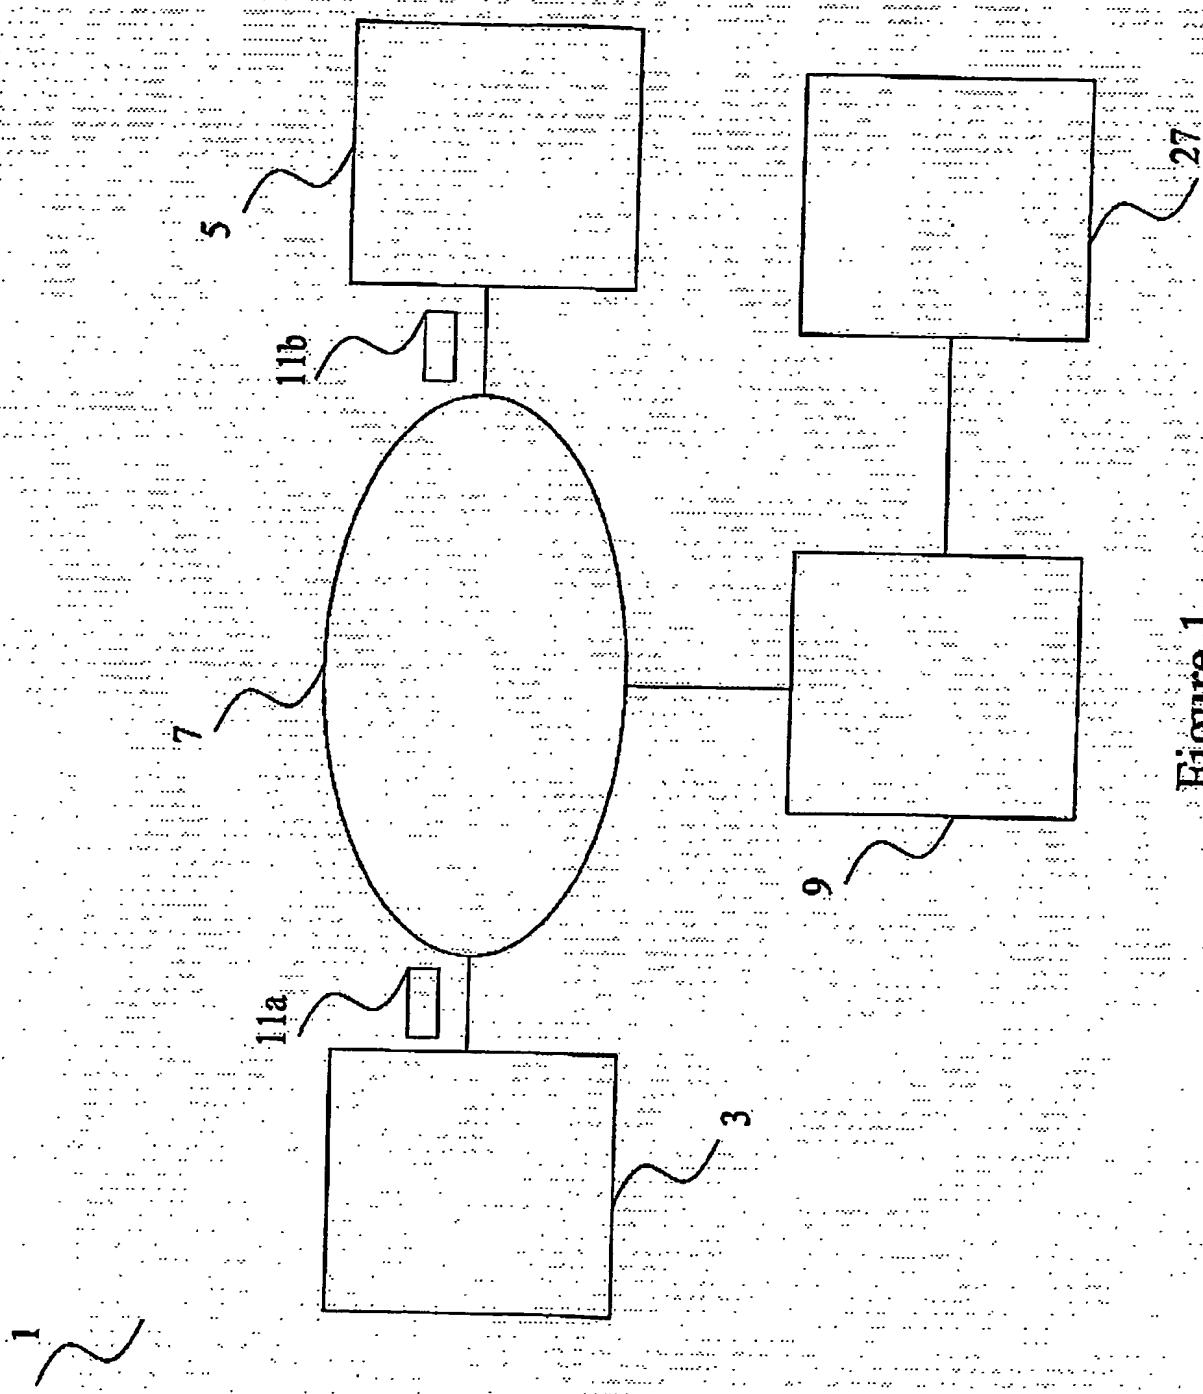
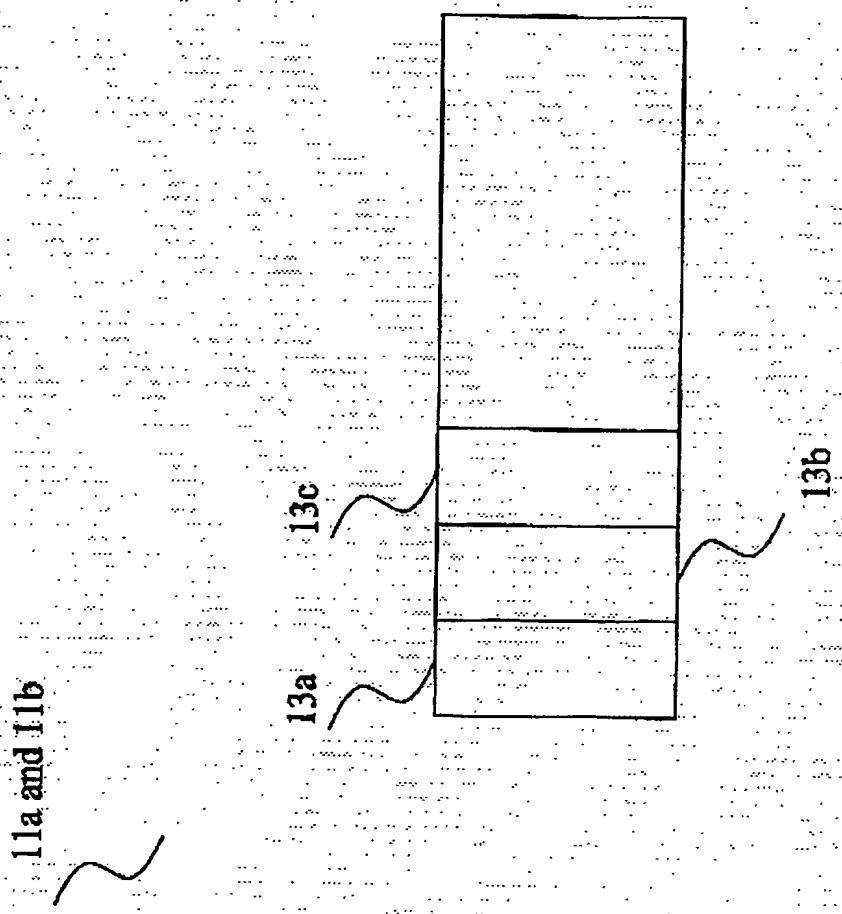


Figure 1

10 '02 THU 16:24 FAX 61 2 99255996

GRIFFITH HACK

021



11a and 11b

Figure 2

02 THU 16:24 FAX 61 2 88255996

GRIFFITH HACK

022

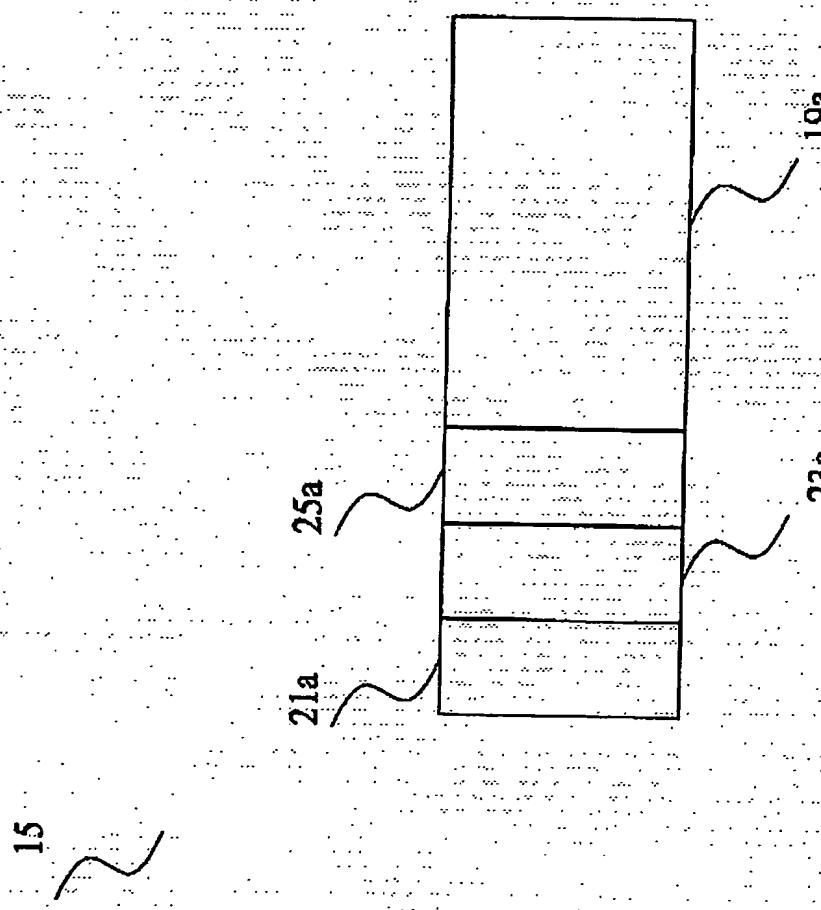


Figure 3

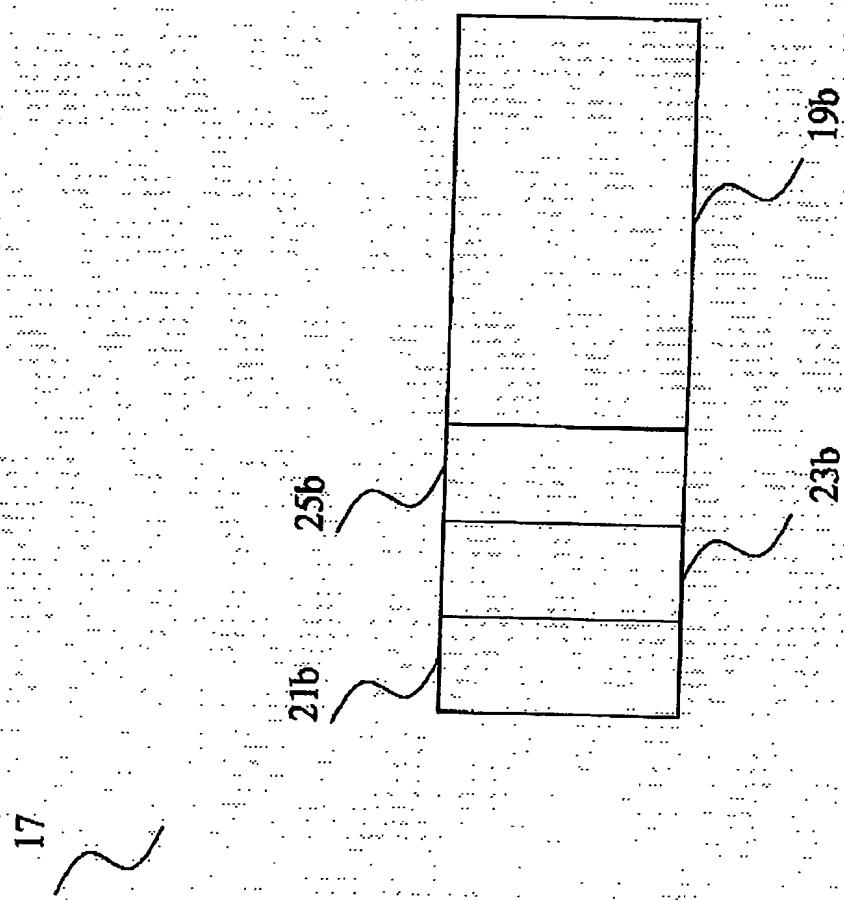


Figure 4

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.